

CONFIDENTIAL

SUN-P7891

UNITED STATES PATENT APPLICATION FOR

A METHOD AND SYSTEM FOR FILESYSTEM MOUNT REDIRECTION

Inventor:

JERRY A. GILLIAM

CHRISTOPHER T. HORNE

SHUDONG ZHOU

Prepared by:

WAGNER, MURABITO & HAO LLP

TWO NORTH MARKET STREET

THIRD FLOOR

SAN JOSE, CALIFORNIA 95113

(408) 938-9060

SUN-P7891/ACM/CWS

A METHOD AND SYSTEM FOR FILESYSTEM MOUNT REDIRECTION

BACKGROUND OF THE INVENTION

FIELD OF THE INVENTION

5 Embodiments of the present invention relate to the field of data storage systems using filesystem services, e.g., with a computer system. More particularly, embodiments of the present invention relate generally to the remounting of a root filesystem, without interrupting filesystem access service, from one device configuration to another in a computer system having mirrored
10 volume disk storage.

RELATED ART

Secondary data storage is an integral part of large data processing systems. A typical data storage system in the past utilized a single, expansive
15 magnetic disk for storing large amounts of data. For single disk memory storage systems, the speed of data transfer to and from the single, large disk is much slower than the processing speed of the Central Processing Unit (CPU) and acts as a data processing bottleneck. In addition, backup capability is needed for improved reliability and serviceability (RAS). In response to these needs,
20 mirrored, redundant arrays of independent disks (RAIDs) have evolved from the single disk storage systems in order to match the speed of secondary storage access with the increasingly faster processing speeds of the CPU and improved RAS. To increase system throughput, the RAID architecture of secondary storage allows for the concurrent access of data from multiple disk drives.

Whether a system has standard single data storage devices or a RAID system of mirrored devices, there will inevitably come a time when a device will be swapped out due to failure, upgrade, etc. To swap out a device that is tied to a filesystem and mount a new one requires that the filesystem be unmounted

5 from the first device and then mounted on the second. This necessitates the termination and restart of the filesystem services. In the case of the root filesystem, this necessitates rebooting the system. In order to reboot, the input/output (I/O) path from the device containing the root filesystem to the drivers for the new device needs to be known. In the case of booting off a

10 mirrored volume device, this information is supplied to the booting firmware by a device that manages the storage devices, known as a Volume Manager. The Volume Manager needs to be initialized as part of the boot process.

Traditionally, the I/O path information has been supplied from the Volume

15 Manager via a forcedload list that is maintained by a system administrator who has the responsibility of maintaining the storage devices. As the complexity of the storage systems has increased, a method has recently become available that allows the Volume Manager to find the drive automatically, independent of the actual path, thus reducing the burden borne by the system administrator.

20 One result of this change is that the path to each of the drivers is no longer a necessarily known quantity. As long as the Volume Manager is running and has I/O available it can dynamically locate the drivers. Therefore, the lack of a defined path would be of no concern. However, this greatly complicates the boot process.

The boot process involves starting with a firmware program such as, in the case of a SUN SPARC platform, for example, an open boot prom (OBP). The OBP firmware takes control of whatever devices to which it has access by using preloaded drivers from firmware. The firmware programs on the OBP are not the same as the software programs that control the hardware when the operating system (e.g., SUN Solaris) is running. Therefore, the OBP and the operating system cannot both be running and controlling the same hardware, as there could be a conflict in instructions that would confuse the hardware. Thus, the OBP must handoff the hardware controls to the operating system kernel without an overlap. One of the main functions of the OBP firmware is to load and initialize the software that the kernel will need to access the storage system, e.g., create a filesystem.

One way to handle this handoff is shown as process 100a of Prior Art Figure 1A. As shown in step 110, the firmware follows a predetermined forcedload list and preloads all the driver modules from disk into memory for the devices that are needed in order to access the root filesystem. These drivers are the ones needed to establish the Volume Manager, in the case of a mirrored system. Then, as shown in step 120, access to I/O is turned off, known as "lights out," at which point the OBP starts running the operating system (OS), but using hardware access via the preloaded driver modules from memory. During this "lights out" period, the root filesystem and the device on which it is mounted (root device) are established by initializing the just loaded drivers, devices, etc. Once the root filesystem is mounted on the root device, the kernel is started and the entire volume of mirrored devices must be established. In the instance of

mirrored volumes, the root device is the Volume Manager. Once the mirrored volume is established, I/O can be regained ("lights on") as exemplified by step 130.

5 Figure 1B is a block diagram showing a mirrored volume in a configuration that may exist following the boot process at the time of "lights on". In this example Filesystem 140 interfaces with the Volume Manager 145 which interfaces with storage devices 150, 155, 160 and 165.

10 If, in order to initialize, the OS needed to bring something in from a disk during "lights out", the OBP must have been provided a path to the disk. Therefore, if any of the driver components that are needed to mount the Volume Manager 145 are missing, the kernel will fail to start after lights out. Some of the driver modules needed by the OBP, for example in the case of RAID storage 15 systems, are obtained from the Volume Manager which builds the forceunload list for the OBP. In the instance where the system is running software in which the Volume Manager finds the drives independent of the actual path, it may not be able to supply the necessary information for the OBP to locate all necessary hardware. In the instance where the system administrator supplies the driver 20 paths, there is always a possibility that there has been an error or omission in the update process.

As upgraded and additional software versions become available and as the mirrored systems become more complex, the burden on the Volume Manager

CONFIDENTIAL

to provide an accurate forceload list becomes great. It would be desirable to reduce this burden on the Volume Manager.

SUMMARY OF THE INVENTION

Accordingly, a need exists for a method and system for allowing the establishment of a filesystem to a storage device with input/output (I/O) remaining available for redirecting the filesystem for access to multiple storage devices, such as mirrored volumes and RAID systems. In other words, a method is needed to eliminate the need, in the case of root filesystem redirection, for completely terminating and subsequently restarting root filesystem services. As a result, the Volume Manager can be freed of the burden of specifying a forceload list because I/O remains active during the initialization of the filesystem.

Specifically, one embodiment of the present invention provides an automated method of establishing a filesystem. The method comprises establishing a first filesystem that interfaces with devices by loading software, including a first set of drivers, into memory and initializing the first set of drivers with the devices.

The first filesystem is then mounted on a root directory that comprises a single storage device. The method allows input/output functionality within the first filesystem and, while input/output functionality is available to the first filesystem, the method accesses the single storage device to obtain software, including a second set of drivers. The method loads the software into the memory and initializes the second set of drivers with the devices to establish a second filesystem. The second filesystem is mounted on a root directory comprising the single storage device and another storage device and the first

filesystem is then rendered inactive. According to one embodiment, the software is determined dynamically.

According to one embodiment, the single storage device and another storage device comprise a mirrored pair of drive devices. In one embodiment, the second filesystem is a volume manager and, according to one embodiment the volume manager directs input/output to and from the mirrored pair of drive devices. According to one embodiment, the devices include drive devices, bus controller devices and bus devices.

10

In one embodiment the method further comprises an application program interfacing with the first filesystem and the second filesystem is established transparently to the application program. The method further comprises the application program interfacing with the second filesystem after the establishment thereof.

15

In one embodiment, a data storage system is disclosed that includes a processor, a plurality of devices and a memory wherein the memory comprises instructions for implementing a method of transparently remounting filesystems.

20

According to one embodiment, an automatic method of establishing a filesystem for accessing a plurality of storage devices is disclosed. The method includes booting from a firmware program that comprises firmware drivers. The booting process includes using the firmware drivers, accessing a first set of drivers from a first storage device and, using the first set of drivers, establishing

25

a first filesystem mounted on a root directory that comprises the first storage device. Once the first filesystem is established, enabled input/output functionality of the first filesystem to the root directory is utilized to access a second set of drivers from the first storage device. The second set of drivers is

5 then used to establish a second filesystem mounted on a root directory. The root directory comprises a number of storage devices including the first storage device. The first filesystem is then rendered inactive.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and form a part of this specification, illustrate embodiments of the invention and, together with the description, serve to explain the principles of the invention:

5

Prior Art Figure 1A is a flow diagram of a method for booting up a system on which a root filesystem is to be remounted in accordance with one embodiment of the conventional art.

10 Prior Art Figure 1B is a block diagram illustrating a configuration of a root filesystem that has been remounted at the time of "lights on" according to one embodiment of the conventional art.

Figure 2 is a logical block diagram of an exemplary computer upon which
15 an embodiment of the present invention may be practiced.

Figure 3A is a block diagram illustrating a filesystem residing on a first device in accordance with an embodiment of the present invention.

20 Figure 3B is a block diagram illustrating a filesystem residing on a second device following redirection in accordance with an embodiment of the present invention.

Figure 4 is a flow diagram illustrating an overview of a computer-implemented process for redirecting a filesystem from a single disk device to a Volume Manager in accordance with one embodiment of the present invention.

5 Figure 5A is a block diagram illustrating a filesystem residing on a disk in accordance with an embodiment of the present invention.

Figure 5B is a block diagram illustrating a having been redirected to a Volume Manager in accordance with an embodiment of the present invention.

10

Figure 6 is a flow diagram of a computer-implemented process for redirecting a filesystem from a disk to a Volume Manager in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

Reference will now be made in detail to the preferred embodiments of the invention, examples of which are illustrated in the accompanying drawings.

While the invention will be described in conjunction with the preferred

5 embodiments, it will be understood that they are not intended to limit the invention to these embodiments. On the contrary, the invention is intended to cover alternatives, modifications and equivalents, which may be included within the spirit and scope of the invention as defined by the appended claims.

Furthermore, in the following detailed description of the present invention,

10 numerous specific details are set forth in order to provide a thorough understanding of the present invention. However, it will be obvious to one of ordinary skill in the art that the present invention may be practiced without these specific details. In other instances, well-known methods, procedures, components, and circuits have not been described in detail so as not to 15 unnecessarily obscure aspects of the present invention.

Notation and Nomenclature

Some portions of the detailed descriptions that follow are presented in terms of procedures, logic blocks, processing, and other symbolic

20 representations of operations on data bits within a computer memory. These descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. In the present application, a procedure, logic block, process, or the like, is conceived to be a self-consistent sequence of steps or 25 instructions leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, although not necessarily, these quantities take the form of electrical or magnetic information capable of being

stored, transferred, combined, compared, and otherwise manipulated in a computer system. It has proven convenient at times, principally for reasons of common usage, to refer to these information as transactions, bits, values, elements, symbols, characters, fragments, pixels, or the like.

5

It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussions, it is appreciated that

10 throughout the present invention, discussions utilizing terms such as

"redirecting," "mounting," "converting," "opening," "flushing," "suspending," "creating," "pointing," "transmitting," "receiving," "closing," "enabling," "generating," or the like, refer to actions and processes of a computer system or similar electronic computing device. The computer system or similar electronic 15 computing device manipulates and transforms data represented as physical (electronic) quantities within the computer system memories, registers or other such information storage, transmission or display devices.

Referring to Figure 2, embodiments of the present invention are comprised 20 of computer-readable and computer-executable instructions that reside, for example, in computer-readable media of an electronic system, such as a host computer system or embedded system. Figure 2 is a block diagram of exemplary embedded components of a computer system 210 upon which embodiments of the present invention may be implemented. Exemplary computer system 210 25 includes an internal address/data bus 220 for communicating information, a central processor 201 coupled with the bus 220 for processing information and

instructions, a volatile memory 202 (e.g., random access memory (RAM), static RAM dynamic RAM, etc.) coupled with the bus 220 for storing information and instructions for the central processor 201, such as instructions for Volume Manager 212, and a non-volatile memory 203 (e.g., read only memory (ROM), 5 programmable ROM, flash memory, EPROM, EEPROM, etc.) coupled to the bus 220 for storing static information such as Open Boot PROM (OBP) firmware 213 and other instructions for processor 201.

With reference still to Figure 2, an optional signal Input/Output (I/O) 10 device 208 is shown. The I/O device 208 is coupled to bus 220 for providing a communication link between computer system 210 and an array network of data storage devices, such as disks. As such, signal I/O device 208 enables the central processor unit 201 to communicate with or monitor other electronic systems blocks that are coupled to the computer system 210. In one 15 embodiment of the present invention, the input and output device 208 is a serial communication port, but could also be any number of well known communication standards and protocols, e.g., Universal Serial Bus (USB), Ethernet, infrared (IR) communication, Bluetooth wireless communication, etc. Instructions and data from the computer system 210 travel through the port and 20 onto an external bus 230 that provides for data transfer between components of the data storage system 204, including between volume manager 212, processor 201 and an array of data disks and associated drives 215.

Figure 3A is a block diagram 300a illustrating a filesystem 310a mounted 25 on a first device 320 in accordance with an embodiment of the present

invention. Filesystem 310a may need to be moved from first device 320, (e.g., a disk), to another device (e.g. second device 330 of Figure 3B). This may be a result of a fault on device 320, the installation of a newer device, a boot-up process or any one of a number of reasons for moving a filesystem from one 5 device to another. In one embodiment, filesystem 310a is mounted on a single disk.

Figure 3B is a block diagram 300b illustrating a second filesystem 310b mounted on a second device 330 or group of devices following redirection as 10 defined below with respect to Figure 4, in accordance with an embodiment of the present invention. Second device 330 may be a disk, a Volume Manager supplying the filesystem data to a set of disks or any of a number of file storage devices known to those skilled in the art of information management.

15 Figures 4 and 6 below are flow diagrams 400 and 600 of computer-implemented processes for redirecting a filesystem from one device to another device in accordance with embodiments of the present invention. Flow diagrams 400 and 600 include processes of the present invention which, in various embodiments, are carried out by a processor and electrical components 20 under the control of computer readable and computer executable instructions. The computer readable and computer executable instructions may reside, for example, in data storage features such as volatile memory 202 and/or non-volatile memory 202 or 203 of Figure 2. However, the computer readable and computer executable instructions may reside in any type of readable storage 25 medium. Although specific steps are disclosed in flow diagram 400, such steps

are exemplary. That is, the present invention is well suited to performing various other steps or variations of the steps recited in Figures 4 and 6. Within the present embodiment, it should be appreciated that the steps of flow diagrams 400 and 600 may be performed by software, by hardware or by any 5 combination of software and hardware.

Figure 4 is a flow diagram 400 of a general overview of the steps in a computer-implemented process for redirecting a filesystem (e.g., 310a of Figure 3A) from a single device (e.g., 320 of Figure 3A) to another device (e.g., 330 of 10 Figure 3B) such as a Volume Manager, in accordance with one embodiment of the present invention.

In step 410, a filesystem is mounted to a single disk. The mounting process involves several steps that include using firmware drivers to obtain the 15 necessary drivers (from the single disk) for mounting the filesystem to the single disk. While the drivers for mounting the filesystem are being loaded to memory, input/output (I/O) activity is suspended so that there is no overlap between the firmware controlling the drivers and the operating filesystem to which the driver control is being transferred. While lights out is in effect, the drivers and 20 hardware for establishing the filesystem are initialized on the single disk. Although only one disk is being used at this point, the single disk may be part of a disk array.

In step 420, according to one embodiment, the I/O activity is reinstated (lights on). The root filesystem 310a is now mounted on the single device 320 and all I/O is available for the following step.

5 Next, in accordance with one embodiment and as illustrated in step 430, the currently mounted filesystem on device 320 is used to load the software objects needed for a Volume Manager 330. With I/O available, any objects determined to be missing from a configuration list may be reported and the path to any missing objects may be determined according to the functionality of the
10 active system. In the case of certain storage systems in which the Volume Manager finds the drives independent of a preconfigured path, the availability of I/O allows the Volume Manager 330 to be loaded. Once loaded, the Volume Manager is initialized. This creates a second filesystem, one including the Volume Manager.

15

In step 440 of Figure 4, the I/O between the second filesystem and device 320 is suspended and the second filesystem is transparently remounted to the Volume Manager 330 so that the mirrored pair system becomes activated. Device 320 is then closed and I/O activity is reinstated to the second filesystem
20 310b via the Volume Manager, in accordance with one embodiment, and the process ends.

Figure 5A is a block diagram 500a illustrating a first filesystem 510a mounted on a disk 520 in accordance with an embodiment of the present
25 invention. Filesystem 510a may need to be moved from first device 520, (e.g., a

disk), to a second device (e.g. Volume Manager 515 of Figure 5B). For example, during a boot up on a computer operating system such as a SUN Microsystems Solaris™ operating environment, having a set of mirrored volumes of disk storage devices, it may be desirable to minimize the paths that

5 the boot software/firmware uses to mount the root filesystem, thus minimizing the time that the input/output (I/O) activity is suspended. One way to do so is, according to an embodiment of the present invention, to mount the root filesystem on a single disk during the boot up. Once the operating system (e.g., Solaris™) is up and running and the I/O is available, it may be remounted on

10 the Volume Manager that, in turn, affords the volume of mirrored devices access to the root filesystem.

Figure 5B is a block diagram illustrating a second filesystem 510b mounted on a Volume Manager 515 following redirection, in accordance with

15 an embodiment of the present invention. Volume Manager 515 manages a mirrored volume of data storage disks 520, 525, 530 and 535, according to one embodiment. The Volume Manager 515 has the responsibility of knowing the location of drivers for the volume of managed disks. In order to know these locations, in some instances, the Volume Manager may need to have I/O activity

20 available.

Figure 6 is a flow diagram of a computer-implemented process 600 for redirecting a filesystem (e.g., filesystem 510a of Figure 5A) from a disk (e.g., disk 520 of Figure 5A) to a Volume Manager (e.g., Volume Manager 515 of

25 Figure 5B) during a boot-up process, in accordance with one embodiment of the

present invention. In step 610, the boot firmware, such as Open Boot PROM (OBP) that is run on the SUN Solaris™ operating system, is initialized. It is appreciated that any combination of hardware and/or software components may be used to realize such a process and OBP is only one example.

5

As illustrated by step 610 of process 600 and in accordance with one embodiment of the present invention, at the time of power on, the OBP is initialized. The initialization of the OBP uses firmware drivers and a defined path to a specified single device 520. Single device 520, according to one 10 embodiment, is one disk of a mirrored volume. From this single device, the OBP obtains the drivers needed to mount first filesystem 510a onto the single device 520.

Whereas, if the filesystem were to be mounted directly to a Volume 15 Manager for managing a system of mirrored pairs, the determination of the drivers and objects necessary for initializing the Volume Manager would be more complex than those needed for mounting to the single device. Also, since the initialization process requires the suspension of I/O so as not to have overlapping systems controlling the same hardware, if drivers or objects were 20 missing or unknown, the initialization would fail without notification of missing objects.

Still referring to Figure 6, in step 620, according to one embodiment, the I/O is suspended (lights out) and the drivers and hardware needed to establish 25 the root filesystem on the single device are initialized. Once the initialization

process is complete, the root filesystem (e.g., filesystem 510a of Figure 5A) is mounted on the single device 520, and the system kernel is up and running.

The I/O is then enabled (lights on).

5 As illustrated by step 630, once the first filesystem is established and the I/O is enabled, the first filesystem may be used to determine and obtain the location of objects and/or devices needed to construct a Volume Manager according to one embodiment. The objects and/or devices may comprise such items as software drivers, bus devices, bus controller devices, bus drivers, or
10 any object or device needed for the complete establishment of the Volume Manager. A character string may furnish the paths to the devices, or they may be determined dynamically by the Volume Manager during its initialization process.

15 In step 640 of process 600, according to one embodiment, the Volume Manager is initialized and I/O activity is suspended between the single device 520 and the root filesystem 510. At this point, a second filesystem 510b has been created that includes the Volume Manager.

20 The root file system 510b is then, in accordance with one embodiment of the present invention and as illustrated in step 650 of Figure 6, transparently transferred from single device 520 to Volume Manager 145. By being transparently transferred, any application program that may be interfacing with the root filesystem as mounted on the single device, at the time of the transfer to
25 the Volume Manager the application program would continue running on the

root filesystem as mounted on the Volume Manager. Since the system kernel is up and running during the transfer, any available system software task may be performed, such as transitioning from the single disk to the mirrored pair. This transition assumes that the data on the two disks of a mirrored pair is identical.

5 At this point, the root filesystem is mounted on the Volume Manager.

In step 660 of Figure 6, the I/O activity is resumed, in accordance with one embodiment of the present invention, thus permitting the second filesystem 510b to now access data from the Volume Manager. The Volume Manager may 10 now direct input/output to and from the mirrored pair of devices. The process is now complete and the flow diagram is exited.

Table 1 below is a glossary of terms as used in ensuing tables of pseudo code that might be used in a Solaris™ operating environment as provided by 15 Sun Microsystems for establishing a new device for mounting a root filesystem. The function commands and terms may include, but are not limited to, those described in the following paragraphs and associated tables.

Table 1

20 I/O: Input/Output. A read from, or a write to, a device.
Block I/O: I/O performed in granularity of a defined block size.
25 filesystem: A service providing directory and file semantics layered on top of a device providing block I/O services.
mount: The act of instantiating a filesystem's services, with I/O operations directed to an underlying device.
30 unmount: The act of terminating a previous mount operation. All filesystem services must cease to complete the unmount.

dev_t: A device reference.

open: Initializing a device instance to accept I/O requests.

5 close: Notifying an open device instance that I/O requests have ceased.

vnode: A file reference. Also called filesystem metadata.

10 inode: A filesystem-specific form of a vnode.

boot: The operation of instantiating the services provided by the system.

15 boot device: The device on which the system is booted. Interchangeable with root device.

root: The filesystem of the boot device.

20 volume: A device composed of multiple physical devices.

25 mirror: A volume composed of devices organized in one or more symmetric subsets, with each subset containing identical data. Redundancy permits recovery in the event of a physical device failure.

SVM: Solaris Volume Manager. The software component of Solaris capable of constructing volumes built on individual physical devices.

Table 2 below illustrates one example of pseudo code that may be used in a Solaris™ operating environment as provided by Sun Microsystems for initializing a new device, (e.g., Volume Manager 515) onto which the root filesystem is to be redirected, as shown in step 630 of Figure 6 in accordance with one embodiment of the present invention.

Table 2

35 open new device

get dev_t used by root filesystem, on currently mounted boot device

40 get dev_t of disk volume device, on which root filesystem is to be remounted

create vnode for new root device

45 open new root device vnode

Table 3 below illustrates one example of pseudo code that may be used in a Solaris™ operating environment as provided by Sun Microsystems for suspending I/O activity between the root filesystem (e.g., filesystem 510 of Figure 5A) and the currently mounted device, (e.g., disk 520 of Figure 5A), as shown in step 640 of Figure 6 in accordance with one embodiment of the present invention.

Table 3

10 suspend I/O activity from filesystem to currently mounted device
 flush any outstanding I/O's to currently mounted device
 suspend root filesystem I/O activity to root device
15

Table 4 below illustrates one example of pseudo code that may be used in a Solaris™ operating environment as provided by Sun Microsystems for transparently transferring filesystem references from the old (currently mounted) device 520 to the new device (e.g., Volume Manager 515), as shown in step 650 of Figure 6 in accordance with one embodiment of the present invention.

Table 4

25 redirect all filesystem references from old device to the new device
 redirect references from original root device to new root device
 convert buf reference of root filesystem inode from original root
 device to new root device
30 convert all cached inodes with an old root dev_t reference to refer to
 the dev_t of the new root device
 point global root filesystem references to new root dev_t and new root
35 vnode

Table 5 below illustrates one example of pseudo code that may be used in a Solaris™ operating environment as provided by Sun Microsystems for unsuspending I/O activity in accordance with the embodiment of step 660 of process 600.

5

Table 5

unsuspend I/O activity, permitting filesystem to now access file data and filesystem metadata from new device.

10 enable root filesystem I/O to new root device
close old device
close old root device
15 release vnode of old root device

Following the redirection of the root filesystem from the originally mounted device, the Volume Manager on which the root filesystem is now mounted is currently providing access for the mirrored volume to the file data and filesystem metadata. In this manner the filesystem has been redirected with no need for complete termination and restart of filesystem services and with I/O availability for determining the driver paths during Volume Manager establishment. The process now exits flow diagram 600.

25

The foregoing descriptions of specific embodiments have been presented for purposes of illustration and description. They are not intended to be exhaustive or to limit the invention to the precise forms disclosed, and many modifications and variations are possible in light of the above teaching. The 30 embodiments were chosen and described in order to best explain the principles

CONFIDENTIAL

of the invention and its practical application, to thereby enable others skilled in the art to best utilize the invention and various embodiments with various modifications as are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the claims appended hereto and

5 their equivalents.